

Configuration Manual

MSc Research Project MSc in Data Analytics

Chethan Sureshbabu StudentID:x21235091

School of Computing National College of Ireland

Supervisor: Mrs Harshani Nagahamulla

National College of Ireland Project Submission Sheet School of Computing



Student Name:	Chethan Sureshbabu
Student ID:	x21235091
Programme:	MSc in Data Analytics
Year:	2023
Module:	MSc Research Project
Supervisor:	Mrs Harshani Nagahamulla
Submission Due Date:	31/01/2024
Project Title:	Configuration Manual
Word Count:	928
Page Count:	14

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

<u>ALL</u> internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Chethan Sureshbabu		
Date:	31st January 2024		
PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:			

Attach a completed copy of this sheet to each project (including multiple copies). **Attach a Moodle submission receipt of the online project submission**, to each project (including multiple copies).

You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Configuration Manual

Chethan Sureshbabu x21235091

1 Introduction

The configuration manual serves as an essential guide, providing a detailed walkthrough for the proper execution of the Multimodal Stress Analysis project. It encompasses stepby-step procedures and crucial information about system requirements, software specifications, and library versions necessary for the successful implementation of the project. Following this manual meticulously ensures a smooth setup and execution, allowing users to delve into the intricate process of analyzing stress across multiple modalities seamlessly.

2 Hardware Specifications

In this section, the hardware specifications of the local machine are described.

2.1 Hardware Specification of the Local Machine

	Item	Value
	OS Manufacturer	Microsoft Corporation
	System Name	LAPTOP-STP96E14
	System Manufacturer	Acer
	System Model	Aspire A715-41G
	System Type	x64-based PC
	System SKU	00000000000000
	Processor	AMD Ryzen 5 3550H with Radeon Vega Mobile Gfx, 2100 Mhz, 4 Core(s), 8
	BIOS Version/Date	INSYDE Corp. V1.01, 05-05-2020
	SMBIOS Version	3.1
	Embedded Controller Version	1.01
	BIOS Mode	UEFI
	BaseBoard Manufacturer	РК
	BaseBoard Product	Azalea_PKS
	BaseBoard Version	V1.01
	Platform Role	Mobile
	Secure Boot State	On
	PCR7 Configuration	Elevation Required to View
	Windows Directory	C:\WINDOWS
	System Directory	C:\WINDOWS\system32
	Boot Device	\Device\HarddiskVolume1
	Locale	United States
	Hardware Abstraction Layer	Version = "10.0.22621.2506"
	User Name	LAPTOP-STP96E14\chethan
	Time Zone	GMT Standard Time
	Installed Physical Memory (RA	8.00 GB
	Total Physical Memory	5.94 GB
	Available Physical Memory	830 MB
	Total Virtual Memory	21.9 GB
	Available Virtual Memory	11.5 GB
	Page File Space	16.0 GB
	Page File	C:\pagefile.sys
ſ		

The above Figure 1 details the Hardware specifications of the local machine to implement or execute the project.

spire A715-41G	
i Device specific	ations
Device name	LAPTOP-STP96E14
Processor	AMD Ryzen 5 3550H with Radeon Vega Mobile Gfx 2.10 GHz
Installed RAM	8.00 GB (5.94 GB usable)
Device ID	232BAE5E-C2F2-4ECA-AF79-3D9959A087A2
Product ID	00327-36230-07351-AAOEM
System type	64-bit operating system, x64-based processor
Windows spec	ifications
Windows spec	ifications Windows 11 Home Single Language
Windows spec Edition	ifications Windows 11 Home Single Language
Windows spec Edition Version	ifications Windows 11 Home Single Language 22H2
Windows spece Edition Version Installed on	ifications Windows 11 Home Single Language 22H2 10-11-2022
Windows spec Edition Version Installed on OS build	ifications Windows 11 Home Single Language 22H2 10-11-2022 22621.2861

Figure 2: Hardware and Device Information

The above Figure 2 shows describes the device specification and Windows specification of the local machine.

3 Software tools required to run/execute the project.

The applications used to implement or complete the project are mentioned below:

- Anaconda navigator 2.5.0
- Jupiter Notebook version 6.5.4

- Python was used as a programming language.
- TensorFlow 2.14.0
- Sklearn 1.3.2

4 Python Libraries were used.

```
1 #importing libraries for text model
   import re
2
   import pickle
3
4 import numpy as np
5 import pandas as pd
6 import seaborn as sns
   from wordcloud import WordCloud
7
8 import matplotlib.pyplot as plt
9
   from nltk.stem import WordNetLemmatizer
10 from sklearn.svm import LinearSVC
11 from sklearn.naive_bayes import BernoulliNB
12 from sklearn.linear_model import LogisticRegression
13 from sklearn.model_selection import train_test_split
14 from sklearn.feature extraction.text import TfidfVectorizer
   from sklearn.metrics import confusion_matrix, classification_report
15
16 import matplotlib.pyplot as plt
17 import time
18 from sklearn.metrics import roc_curve, auc
19 from joblib import dump
```

Figure 3: Importing libraries for text model.

```
1 #importing libraries for audio model
         import os
import librosa
          import pandas as pd
 4
 5
          import numpy as np
 6
         from sklearn.model_selection import train_test_split
            import os
 7
 8
         import pandas as pd
 9
            import numpy as np
           from sklearn.model_selection import train_test_split
0
            from keras.utils import to_categorical
.1
.2
            from sklearn.preprocessing import StandardScaler, LabelEncoder
3
            from keras.models import Sequential
           from keras.layers import Dense, Dropout
from sklearn.metrics import classification_report, accuracy_score
4
5
.6
            from sklearn.model_selection import train_test_split
            import matplotlib.pyplot as plt
 8
            import seaborn as sns
            from sklearn.metrics import confusion_matrix, roc_curve, auc
9
Image: Propert of the second secon
```

Figure 3: Importing libraries for an audio model.

```
1 #importing libraries for image model
  import os
3 import numpy as np
4 import pandas as pd
5 from sklearn.model_selection import train_test_split
6 from sklearn preprocessing import LabelEncoder, MinMaxScaler
  from sklearn.metrics import classification_report, confusion_matrix, roc_curve, auc
8 from tensorflow.keras.preprocessing.image import load_img, img_to_array
  from tensorflow.keras.models import Sequential
9
.0 from tensorflow.keras.layers import Conv2D, MaxPooling2D, Flatten, Dense, Dropout
  from tensorflow.keras.callbacks import EarlyStopping
  from tensorflow.keras.optimizers import Adam
  from tensorflow.keras.applications import EfficientNetB0
4 from efficientnet.keras import center_crop_and_resize
  import seaborn as sns
5
.6 import matplotlib.pyplot as plt
.7 from joblib import dump
```

Figure 4: Importing libraries for an image model.

```
1 #importing libraries for late fusion model
2 import numpy as np
3 from joblib import load, dump
4 from sklearn.model selection import train test split
5 from sklearn.preprocessing import StandardScaler
6 from sklearn.linear_model import LogisticRegression
  from sklearn.metrics import classification report, confusion matrix
7
8 from sklearn.feature extraction.text import TfidfVectorizer
9 from sklearn.metrics import matthews corrcoef
Ø from sklearn.metrics import roc_auc_score
1 from sklearn.metrics import auc
2 from sklearn.calibration import calibration curve
3
  from sklearn.metrics import log_loss
4 from sklearn.metrics import brier_score_loss
```

Figure 5: Importing libraries for a Late Fused model.

The above figure shows the description of the Python libraries we have used while implementing Text, Audio, Image, and Late Fusion models.

6 Datasets used in the project.

Dataset 1 - The sentiment140 dataset is used for the Text Model, It contains 1,600,000 tweets extracted using the Twitter API.

Dataset 2 - Facial Expressions Training Data is used for the Image Model implementation, This data set is based on Affect Net-HQ, Affect Net is a large database of faces labeled by "affects"

Dataset 3 - RAVDESS Emotional speech audio data set is used for implementing the Audio model. This portion of the RAVDESS contains 1440 files The RAVDESS contains 24 professional actors (12 female, 12 male), vocalizing two lexically matched statements in a neutral North American accent. Speech emotions include calm, happy, sad, angry, fearful, surprised, and disgusted expressions.

So, in this section, we have discussed the three data sets we have used for the Multimodal Stress Analysis project.

7 Code Snippets

In this section, important code snippets from the Text, Audio, Image, and Late Fusion models we have implemented will be discussed.

7.1 Text Model Implementation

-	1 2 3	<pre># Display unique values in the 'sentiment' column unique_sentiments = dataset['sentiment'].unique() print("Unique Sentiments:", unique_sentiments)</pre>						
	<pre>4 5 # Adjust the mapping based on your sentiment values 6 sentiment_mapping = {0: 'high stressed', 4: 'low stressed'} 7 dataset['stress_level'] = dataset['sentiment'].map(sentiment_mapping)</pre>							
	<pre>8 9 # Checking the class counts 10 class_counts = dataset['stress_level'].value_counts() 11</pre>							
	12 13	print("(print(c	Class Coun lass_count	ts:") s)				
:	<pre>Unique Sentiments: [0 4] Class Counts: stress_level low stressed 800000 high stressed 799999 Name: count, dtype: int64 1 # Storing data in Lists 2 text, stress_levels = list(dataset['text']), list(dataset['stress_level']) 3</pre>							
:	1	dataset	.head()					
÷	S	entiment	ids	date	flag	user	text	stress_level
	0	0	1467810672	Mon Apr 06 22:19:49 PDT 2009	NO_QUERY	scotthamilton	is upset that he can't update his Facebook by \ldots	high stressed
	1	0	1467810917	Mon Apr 06 22:19:53 PDT 2009	NO_QUERY	mattycus	@Kenichan I dived many times for the ball. Man	high stressed
	2	0	1467811184	Mon Apr 06 22:19:57 PDT 2009	NO_QUERY	ElleCTF	my whole body feels itchy and like its on fire	
								high stressed
	3	0	1467811193	Mon Apr 06 22:19:57 PDT 2009	NO_QUERY	Karoli	@nationwideclass no, it's not behaving at all	high stressed high stressed

Fig 6: Represents the Mapping function.

After loading the text data set we applied the mapping function to the data to add the stress levels column in the dataset.

```
1 import time
 2 t = time.time()
 3 processedtext = preprocess(text)
4 print(f'Text Preprocessing complete.')
 5 print(f'Time Taken: {round(time.time()-t)} seconds')
Text Preprocessing complete.
Time Taken: 308 seconds
 1 # Splitting the Data
 2 X_train, X_test, y_train, y_test = train_test_split(processedtext, stress_levels, test_size=0.05, random_state=0)
 3
 1 # Create the folder if it doesn't exist
2 save_path = 'C:\\Users\\chethan\\OneDrive\\Desktop\\Text Test Data'
 3 os.makedirs(save_path, exist_ok=True)
 1 # Combine the text and labels for the selected test samples
2 selected_test_data = list(zip(X_test, y_test))
 1 # TF-IDF Vectorizer
 2 vectoriser = TfidfVectorizer(ngram_range=(1, 2), max_features=500000)
 3 vectoriser.fit(X_train)
 1 # Transforming the data
 2 X_train = vectoriser.transform(X_train)
 3 X_test = vectoriser.transform(X_test)
 1 # Evaluate Model Function
 2 def model_Evaluate(model):
        y_pred = model.predict(X_test)
print(classification_report(y_test, y_pred))
cf_matrix = confusion_matrix(y_test, y_pred)
 3
 4
  5
```

Figure 7: Important preprocessing steps

In figure 7 we can see the most important preprocessing steps we have we have done before building the model.

<pre>1 # Logistic Regression ModeL 2 LRmodel = LogisticRegression(C=2, max_iter=1000, n_job 3 LRmodel.fit(X_train, y_train) 4 model_Evaluate(LRmodel)</pre>					obs=-1)
	precision	recall	f1-score	support	
high stresse	d 0.83	0.82	0.83	39986	
low stresse	d 0.82	0.83	0.83	40014	
accuracy	y		0.83	80000	
macro av	g 0.83	0.83	0.83	80000	
weighted av	g 0.83	0.83	0.83	80000	

Figure 8: Model building

In the above code snippet we can see the code for implementing the Logistic Regression model and also the classification report of the model.



Figure 9: Saving the Logistic Regression model

Here, we have saved the model to use in the future for the Late Fusion Model

7.2 Audio Model Implementation

```
# Function to load RAVDESS dataset
def load_ravdess_dataset(dataset_path):
    data = {'path': [], 'emotion': []}
    # Iterate over actor folders
    for actor_folder in os.listdir(dataset_path):
        actor_path = os.path.join(dataset_path, actor_folder)
        if os.path.isdir(actor_path):
            # Iterate over audio files in each actor folder
            for audio_file in os.listdir(actor_path):
               audio_path = os.path.join(actor_path, audio_file)
                # Extract emotion from the file name (assuming file names are structured)
                emotion = int(audio_file.split('-')[2])
                # Append data to the dictionary
                data['path'].append(audio_path)
               data['emotion'].append(emotion)
    # Create a DataFrame from the dictionary
    df = pd.DataFrame(data)
    return df
# path to the dataset
dataset_path = r'C:\\Users\\chethan\\Downloads\\RIC\\audio_speech_actors_01-24'
# Loading the dataset
ravdess_dataset = load_ravdess_dataset(dataset_path)
```

Figure 10: Function to load the Audio dataset.

Using the above function we have successfully loaded the Audio dataset.

```
1
        # Mapping emotion labels to stress levels
   2
        stress_mapping = {
   3
            1: 'low_stress',
             2: 'low_stress',
3: 'low_stress',
   4
   5
             4: 'low_stress'
   6
             5: 'high_stress',
   7
             6: 'high_stress',
7: 'high_stress',
   8
   9
             8: 'high_stress',
  10
 11 }
  12
  13 ravdess_dataset['stress_level'] = ravdess_dataset['emotion'].map(stress_mapping)
 14
   1 # Displaying the loaded dataset
   2 print(ravdess_dataset.head())
                                                                                    path emotion stress_level

      0 C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
      1 low_stress

      1 C:\\Users\\chethan\\Downloads\\RIC\\audio spee...
      1 low_stress

1 C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
                                                                                                          1 low_stress

      1
      C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
      1
      low_stress

      2
      C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
      1
      low_stress

      3
      C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
      1
      low_stress

      4
      C:\\Users\\chethan\\Downloads\\RIC\\audio_spee...
      2
      low_stress
```

Figure 11: Mapping the stress levels.

Then using the above function we have mapped the stress levels in the dataset.

```
1 # Function to extract features from audio files
2
  def extract_features(file_path):
3
      try:
4
          # Loading audio file with librosa
          audio, sample_rate = librosa.load(file_path, res_type='kaiser_fast')
5
6
7
          # Extracting features using librosa
          mfccs = np.mean(librosa.feature.mfcc(y=audio, sr=sample_rate, n_mfcc=13), axis=1)
8
9
          chroma = np.mean(librosa.feature.chroma_stft(y=audio, sr=sample_rate), axis=1)
          mel = np.mean(librosa.feature.melspectrogram(y=audio, sr=sample rate), axis=1)
0
          contrast = np.mean(librosa.feature.spectral_contrast(y=audio, sr=sample_rate), axis=1)
1
          tonnetz = np.mean(librosa.feature.tonnetz(y=audio, sr=sample_rate), axis=1)
2
.3
4
          # Combining all features into a single array
.5
          feature_array = np.concatenate((mfccs, chroma, mel, contrast, tonnetz))
6
7
      except Exception as e:
8
          print(f"Error encountered while parsing file: {file_path}. Error: {e}")
9
          return None
0
1
      return feature_array
1 # Extracting features from each audio file and create a DataFrame
2
  !pip install resampy
```

ravdess_dataset['features'] = ravdess_dataset['path'].apply(extract_features)

3

4

Figure 12: Function to extract the features.

Using the above function we can extract the audio features from the audio data set.

```
1 # Converting features to a list and labels to a NumPy array
2 X = np.array(ravdess_dataset['features'].tolist())
3 y = np.array(ravdess_dataset['stress_level'])
1 # Encoding the target labels
2 le = LabelEncoder()
3 y_encoded = le.fit_transform(y)
4 y_categorical = to_categorical(y_encoded)
1 # Splitting the dataset into training and testing sets
2 X_train, X_test, y_train, y_test = train_test_split(X, y_categorical, test_size=0.2, random_state=42)
3 
1 # Standardizing the features
2 scaler = StandardScaler()
3 X_train_scaled = scaler.fit_transform(X_train)
4 X_test_scaled = scaler.transform(X_test)
```

Figure 13: The important preprocessing steps before implementing the model.



Figure 14: CNN Model implementation.

The above code is used for model implementation.

```
1 # Making predictions on the test set
 2 y pred = model.predict(X test scaled)
9/9 [=====] - 0s 5ms/step
 1 # Convertting predictions to class labels
 2 y_pred_labels = le.inverse_transform(np.argmax(y_pred, axis=1))
 3 y_test_labels = le.inverse_transform(np.argmax(y_test, axis=1))
 1 # Evaluate the model
 2 print("Classification Report:\n", classification_report(y_test_labels, y_pred_labels))
 3 print("Accuracy:", accuracy_score(y_test_labels, y_pred_labels))
Classification Report:
             precision recall f1-score support
                          0.87
                                   0.85
 high_stress
                0.83
                                             151
 low_stress
                0.85
                          0.80
                                   0.82
                                             137
                                   0.84
                                             288
   accuracy
                                0.84
                0.84 0.84
0.84 0.84
                                              288
  macro avg
weighted avg
                                   0.84
                                              288
Accuracy: 0.836805555555556
```

Figure 15: Model prediction on test data.

```
1 from joblib import dump
2
3 # Save the audio model
4 save_path = 'C:/Users/chethan/Downloads/RIC/audio_model_saved.joblib'
5 dump(model, save_path)
6
7 print(f"Audio model saved successfully at: {save_path}")
```

Audio model saved successfully at: C:/Users/chethan/Downloads/RIC/audio_model_saved.joblib

Figure 16: saving the model.

7.3 Image Model Implementation

```
2 path = 'C:\\Users\\chethan\\Downloads\\RIC\\Image dataset\\'
3 file = os.path.join(path, 'labels.csv')
4 # Load the dataset
5 emotion_dataset = pd.read_csv(file)
6 emotion_dataset.head()
  Unnamed: 0
                              pth
                                    label
                                          relFCs
0
           0 anger/image0000006.jpg surprise 0.873142
1
           1 anger/image0000060.jpg
                                    anger 0.852311
           2 anger/image0000061.jpg
                                   anger 0.800957
2
3
           3 anger/image0000066.jpg
                                   disgust 0.843079
           4 anger/image0000106.jpg
4
                                    anger 0.849108
```

Figure 17: Loading the image dataset

```
1 # Map existing labels to stress levels
 2 stress_level_mapping = {
3 'anger': 'high_stressed'.
          'contempt': 'low_stressed',
'disgust': 'high_stressed',
'fear': 'high_stressed',
 4
 5
  6
          'happy': 'low stressed',
  7
          'neutral': 'low_stressed',
 8
          'sad': 'high_stressed',
'surprise': 'low_stressed'
 9
10
11 }
  1 # Add a new column 'stress_level' based on the mapping
 2 emotion_dataset['stress_level'] = emotion_dataset['label'].map(stress_level_mapping)
  3
```

```
1 # Adding the 'stress_level' column in your DataFrame
2 high_stressed_count = emotion_dataset[emotion_dataset['stress_level'] == 'high_stressed'].shape[0]
3 low_stressed_count = emotion_dataset[emotion_dataset['stress_level'] == 'low_stressed'].shape[0]
4 print("Number of High Stressed Instances:", high_stressed_count)
5 print("Number of Low Stressed Instances:", low_stressed_count)
6
```

```
Number of High Stressed Instances: 13118
Number of Low Stressed Instances: 15057
```

Figure 18: Mapping the stress levels.

1 2 3 4	<pre># Min-Max Scaling scaler = MinMaxScaler() emotion_dataset[['relFCs']] = scaler.fit_transform(emotion_dataset[['relFCs']])</pre>
1 2 3 4	<pre># Train-test split X_train_features, X_test_features, y_train, y_test = train_test_split(emotion_dataset[['relFCs', 'pth']], emotion_dataset['stress_level'], test_size=0.2, random_state=42)</pre>
1 2 3 4 5 6 7 8 9 10	<pre># Data augmentation using ImageDataGenerator train_datagen = ImageDataGenerator(rotation_range=20, width_shift_range=0.2, height_shift_range=0.2, shear_range=0.2, zoom_range=0.2, horizontal_flip=True, fill_mode='nearest')</pre>
1 2 3 4	<pre># Reshape the data for RGB images image_height = 96 image_width = 96 num_channels = 3 # For RGB images</pre>
1 2 3 4 5 6 7 8 9	<pre># Load and preprocess the images for training with data augmentation X_train_images = [] for image_path in X_train_features['pth']: img = load_img(os.path.join(path, image_path), target_size=(image_height, image_width)) img_array = img_to_array(img) img_array = train_datagen.random_transform(img_array) X_train_images = np.array(X_train_images) </pre>
1 2 3 4 5 6 7 8	<pre># Load and preprocess the images for testing X_test_images = [] for image_path in X_test_features['pth']: img = load_img(os.path.join(path, image_path), target_size=(image_height, image_width)) img_array = img_to_array(img) X_test_images.append(img_array) X_test_images = np.array(X_test_images)</pre>

Figure 19: Important preprocessing steps before implementing the model.

1	# EfficientNetB0 model				
2	def huild efficientnet model(innut shane=(image height image width num channels) num classes=1).				
2	bace model = EfficientNatDQ(input charactionut charaction include ton-Ealso weekstart)				
-	base_model = Efficienceco(input_snape-input_snape, include_top-raise, weights- imagenet)				
4					
5	<pre>model = Sequential()</pre>				
6	model.add(base_model)				
7	model.add(Flatten())				
8	<pre>model.add(Dense(128, activation='relu'))</pre>				
9	model.add(Dropout(0.5))				
10	model.add/Dense(num_classes_activation='sigmoid'))				
11	# #Compile the model				
12	model complete the model				
12	model.compile(optimizer=Adam(learning_rate=0.0001), ioss= binary_crossentropy , metrics=[accuracy])				
13					
14	return model				
1	# Initialize the model				
2	<pre>model = build efficientnet model()</pre>				
4					
1	# Define early stopping to prevent overfitting				
2	early_stopping = EarlyStopping(monitor='val_loss', patience=3, restore_best_weights=True)				
3					
1	# Encode the taraet variable 'stress level'				
2	label encoder = LabelEncoder()				
2	v train encoded - label encoden fit transform(v train)				
	y tort encoded = label_encoder.th_enarcon m(y_c) all)				
4	<pre>y_test_encoded = idet_encoder.transform(y_test)</pre>				
1	# Train the model				
2	history = model.fit(X_train_images, y_train_encoded, epochs=10, batch_size=32,				
3	validation split=0.2, callbacks=[early stopping])				
4					
Epo	ch 1/10				
559,	/559 [========================] - 846s 1s/step - loss: 0.6692 - accuracy: 0.6289 - val_loss: 0.5827 - val accuracy: 0.				
7018					
Enor	ch 2/10				
500	(EEG [
ופכנ					
/31.	1				

Figure 20: EfficientnetBO Model implementation.

```
1 # Evaluate the model
2 test_loss, test_acc = model.evaluate(X_test_images, y_test_encoded)
3 print(f'Test Accuracy: {test_acc}')
4
```

```
175/175 [===============] - 56s 317ms/step - loss: 0.5047 - accuracy: 0.7591
Test Accuracy: 0.7590900659561157
```

```
1 # Make predictions
2 y_pred = model.predict(X_test_images)
3 y_pred_binary = (y_pred > 0.5).astype(int)
```

```
175/175 [======] - 59s 316ms/step
```

```
1 # Classification report
   print(classification_report(y_test_encoded, y_pred_binary))
 2
 3
             precision recall f1-score support
                 0.74 0.74 0.74
0.77 0.78 0.77
                                             2607
          0
                 0.77
                           0.78
                                     0.77
                                               2976
          1
                                     0.76
   accuracy
                                             5583
  macro avg 0.76 0.76 0.76 5583
ghted avg 0.76 0.76 0.76 5583
weighted avg
```

Figure 21: Model prediction on test data.

```
: 1 from joblib import dump
2
3 save_path_efficientnet = 'C:/Users/chethan/Downloads/RIC/efficientnet_model.joblib'
4
5 # Save the EfficientNet model
6 dump(model, save_path_efficientnet)
7
8 print(f"EfficientNet model saved successfully at: {save_path_efficientnet}")
```

EfficientNet model saved successfully at: C:/Users/chethan/Downloads/RIC/efficientnet_model.joblib

Figure 22: Saving the model for using later for the late fusion model.

7.4 Late Fusion Model Implementation

```
10 # Load the saved models
11 audio_model = load('C:/Users/chethan/Downloads/RIC/audio_model_saved.joblib')
12 image_model = load('C:/Users/chethan/Downloads/RIC/efficientnet_model.joblib')
13 text_model = load('C:/Users/chethan/Downloads/RIC/logistic_regression_model.joblib')
14 tfidf_vectorizer = load('C:/Users/chethan/Downloads/RIC/tfidf_vectorizer.joblib')
```

Figure 23: Loading the Text, Audio, and image Models.

```
10 # Load the saved models
1 audio_model = load('C:/Users/chethan/Downloads/RIC/audio_model_saved.joblib')
12 image_model = load('C:/Users/chethan/Downloads/RIC/efficientnet_model.joblib')
13 text_model = load('C:/Users/chethan/Downloads/RIC/logistic_regression_model.joblib')
14 tfidf_vectorizer = load('C:/Users/chethan/Downloads/RIC/tfidf_vectorizer.joblib')
16 # test data for audio, image, and text
17 audio_test_data = X_test_audio
18 video_test_data = X_test_images
19 text_test_data = X_test_text
20
21 # Obtain predictions from individual models
22 audio_predictions = audio_model.predict(audio_test_data)
23 image_predictions = image_model.predict(video_test_data)
24 text_predictions_proba = text_model.predict_proba(tfidf_vectorizer.transform(text_test_data))[:, 1]
26 # Ensuring all predictions have the same number of samples
27 min_samples = min(len(audio_predictions), len(image_predictions), len(text_predictions_proba))
28 audio_predictions = audio_predictions[:min_samples]
29 image_predictions = image_predictions[:min_samples]
30 text_predictions_proba = text_predictions_proba[:min_samples]
31
32 # Concatenate features
33 all_features = np.concatenate((audio_predictions[:, 0].reshape(-1, 1), image_predictions, text_predictions_proba.reshape(-1,
34
35 # Standardizing the features
36 scaler = StandardScaler()
37 X_test_scaled_late = scaler.fit_transform(all_features)
38
39 y_test_binary = np.where(y_test_combined[:, 0] == 'high_stressed', 1, 0)
40 y_test_binary = y_test_binary[:min_samples]
41
42 # Initialize late fusion model
43 late_fusion_model = LogisticRegression()
44
45 late_fusion_model.fit(X_test_scaled_late, y_test_binary)
46
47 # Obtain predictions from the late fusion model
48 late_fusion_predictions_proba = late_fusion_model.predict_proba(X_test_scaled_late)[:, 1]
49
50 # Adjust the decision threshold based on precision-recall curve
51 precision, recall, thresholds = precision_recall_curve(y_test_binary, late_fusion_predictions_proba)
53 # Finding the threshold that maximizes F1-score or any other preferred metric
54 best_threshold = thresholds[np.argmax(2 * precision * recall / (precision + recall))]
```

55 late_fusion_predictions = (late_fusion_predictions_proba >= best_threshold).astype(int)

Figure 23: Late Fusion model implementation.

```
57 # Evaluating performance
58 print("Binary Classification Model Performance:")
59 print(classification_report(y_test_binary, late_fusion_predictions))
60 print("Confusion Matrix:")
61 print(confusion_matrix(y_test_binary, late_fusion_predictions))
62
    9/9 [======] - 0s 4ms/step
175/175 [==================] - 70s 383ms/step
Binary Classification Model Performance:
            precision recall f1-score support
          0
                 0.86
                      0.75
                                   0.80
                                             159
                 0.73
                          0.85
                                   0.79
                                             129
          1
                                   0.80
                                             288
   accuracy
                 0.80
                         0.80
                                   0.79
                                             288
  macro avg
weighted avg
                0.80
                         0.80
                                   0.80
                                             288
Confusion Matrix:
[[119 40]
 [ 19 110]]
```

Figure 24: Classification report of late fusion model.